

What Makes a Good Molecular-Scale Computer Device?

Seth Copen Goldstein Dan Rosewater

September 26, 2002

CMU-CS-02-181

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

The lithographically-produced CMOS transistor has been the key technology that has enabled the information revolution. However, in the near future the limitations, both technical and economic, introduced by lithographic fabrication may inhibit further decreases in feature size. Chemically assembled electronic nanotechnology (CAEN) is a promising alternative to CMOS for constructing circuits with device sizes in the tens of nanometers, far smaller than is thought possible using lithography. In this paper we examine and contrast the constraints imposed by lithographic versus CAEN fabrication; the key limitation is that three-terminal devices, such as transistors, will be impractical at the nanoscale. We demonstrate that these constraints can be satisfied by outlining an architecture that uses only two-terminal CAEN devices to compute without transistors. One crucial requirement of this design circuit is that it be able to restore signals to a reference state without transistors. We present preliminary results for a molecular latch, constructed from molecular resonant tunneling diodes (RTDs) that can perform signal restoration, I/O isolation, and voltage buffering without transistors at the nanoscale.

| Report Documentation Page | | | Form Approved OMB No. 0704-0188 | | |
|--|------------------------------------|-------------------------------------|--|---|---------------------------------|
| Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. | | | | | |
| 1. REPORT DATE 26 SEP 2002 | | 2. REPORT TYPE | | 3. DATES COVERED 00-00-2002 to 00-00-2002 | |
| 4. TITLE AND SUBTITLE What Makes a Good Molecular-Scale Computer Device? | | | 5a. CONTRACT NUMBER | | |
| | | | 5b. GRANT NUMBER | | |
| | | | 5c. PROGRAM ELEMENT NUMBER | | |
| 6. AUTHOR(S) | | | 5d. PROJECT NUMBER | | |
| | | | 5e. TASK NUMBER | | |
| | | | 5f. WORK UNIT NUMBER | | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University, School of Computer Science, Pittsburgh, PA, 15213 | | | 8. PERFORMING ORGANIZATION REPORT NUMBER | | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | | 10. SPONSOR/MONITOR'S ACRONYM(S) | | |
| | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited | | | | | |
| 13. SUPPLEMENTARY NOTES The original document contains color images. | | | | | |
| 14. ABSTRACT | | | | | |
| 15. SUBJECT TERMS | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES 15 | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT unclassified | b. ABSTRACT unclassified | c. THIS PAGE unclassified | | | |

Keywords: Chemically assembled electronic nanotechnology, two-terminal devices, molecular latch, resonant tunneling diodes, clocking scheme

1 Introduction

In 1985, Keyes argued that the silicon transistor would not be easily replaced as the building block in digital circuits [1]. Although faster switching technologies have been introduced, none possess the wide array of desirable features to match silicon: gain, voltage restoration, noise reduction, tolerance to variability in the manufacturing process, isolation of inputs from outputs, excellent fan-out, the ability to invert a signal, tolerance to the high demands of a system environment, and a huge body of research and financial commitment to their use and manufacture. Since that time, advances in CMOS technology have increased the speed and decreased the size of transistors at a phenomenal rate; recent research has produced 30-nm gate CMOS transistors that switch in less than 2 ps [2].

Impressive as these results have been, in the near future further increases in the performance of silicon-based, lithographically manufactured transistors will be difficult to achieve. Never before has there been so much doubt in the industry about how three-generation-removed advances will be accomplished [3, 4]. There are several major reasons for this. First, the small line widths necessary for next-generation lithography require the use of increasingly shorter-wavelength light, which introduces a host of problems that are currently being addressed [5]. In addition, as the number of atoms that constitutes a device decreases, manufacturing variability of even a few atom widths can become significant and lead to defects.

More important, however, is the economic barrier to commercial nanometer-scale lithography. New fabrication facilities orders of magnitude more expensive than present ones will be needed to produce chips with the required densities while maintaining acceptably low defect rates. The increasing cost of chip masks, which must be manufactured to single-atom tolerances, precludes commercially viable development of new chips except for the highest-volume ICs. It is entirely possible that further reduction of transistor size will be halted by economic rather than technological factors.

This economic downside is a direct consequence of the precision required for deep-submicron lithographic fabrication. Using lithography, construction of devices and their assembly into circuits occur at the same time. Keyes pointed out that this is a great advantage of silicon in that mass fabrication of transistors is extremely inexpensive per transistor. However, as Keyes also highlighted, it produces a set of constraints: each element in the system must be highly reliable, and devices cannot be tested, adjusted, or repaired after manufacture. These constraints force the design of a custom circuit to be tightly integrated with manufacture, since no additional complexity can be introduced into the circuit afterwards. Lithography is perfectly tuned to deal with these constraints, but at small scales the precision required to obey these constraints becomes the obstacle to further improvement. Any technology hoping to displace lithographically produced CMOS integrated circuits must overcome such obstacles while sacrificing neither performance nor the low per-unit cost made possible by mass production.

2 Electronic Nanotechnology

Chemically assembled electronic nanotechnology (CAEN) holds promise as a technology to overcome these obstacles. CAEN takes advantage of chemical synthesis techniques to construct molecular-sized circuit elements such as resistors, transistors, diodes [6], resonant tunneling diodes (RTDs) [7,

8], and reconfigurable switches [9, 10]. Chemical synthesis can produce enormous quantities (moles) of identical devices cheaply, or devices can be grown in situ. In all cases, the fabricated devices are only a few nanometers in size and exploit quantum mechanical properties to control voltage and current levels across the terminals of the device. Most of these molecules are smaller than the physical feature-size limit that can be produced using silicon. Additionally, tiny amounts of current are needed to operate these devices, which should result in incredibly low power consumption in devices created from them.

Molecular-scale devices cannot by themselves overcome the constraints that will prove limiting for CMOS unless a suitable, low-cost manufacturing technique can be devised for composing them into circuits. For devices created separately from the circuit assembly process, localization and connection of the devices using a lithographic-like process will be very difficult, and other means must be sought. Connecting the nanoscale components one at a time will be prohibitively expensive due to their small size and the large number of components to be connected. Instead, economic viability requires that circuits be created and connected through self-assembly and self-alignment. Several groups have recently demonstrated CAEN devices that are self-assembled or self-aligned, or both [10, 11, 12, 13]. Advances have also been made in creating wires out of single-wall carbon nanotubes and aligning them on a silicon substrate [14, 15]. More recently, selective destruction has been used to create desired carbon nanotube structures [16]. In addition, metallic nanowires, which can scale down to 5nm and can include embedded devices or device coatings, have been fabricated [17, 18].

Self-assembly of devices to form circuits imposes an entirely different set of constraints from lithography. Because self-assembly will be an imprecise, thermodynamically controlled process, we can no longer expect to produce working devices in specified locations with the high reliability of CMOS lithography. This implies that the devices must be testable, adjustable, and possibly repairable after manufacture, and also that some defects must be tolerated. Moreover, self-assembly can produce only a limited range of circuit complexity while remaining manageable. Thus the manufacturing process must be decoupled from the production of the final circuit. Any deterministic aperiodic circuit will have to be created through reconfiguration after the device is fabricated.

Reconfigurability will therefore be a key component of any successful nanoscale architecture. The regular structures amenable to creation via self-assembly are ideally suited to creating a reconfigurable computing fabric. Reconfigurable computing uses large arrays of uniform, programmable logic elements and interconnects, which can be configured using electrical signals, to perform required computations. Reconfigurability has been used for years to implement new circuits after manufacture, as in commercial and research Field Programmable Gate Arrays [19, 20, 21, 22]. It has also been used to incorporate defective circuit elements into a working, high-performance computer, the Teramac [23], by discovering the defective circuit elements after fabrication.

Molecules that can act as reconfigurable switches in series with diodes have already been discovered [9, 10]. These switch/diode pairs have a great advantage over their CMOS counterparts: configuration information is stored in the device itself, rather than in separate static RAM cells, greatly decreasing the area required for each circuit element. A single CAEN switch-diode pair will require 800nm^2 as opposed to $100,000\text{nm}^2$ for a single laid-out CMOS transistor [24]. A simple CMOS logic gate or static memory cell requires several transistors, separate p- and n-wells, and wiring interconnects, resulting in a factor of more than one hundred difference in density between CAEN and CMOS. While one switch/diode does not implement the same functionality as a

transistor, clearly the size differences give CAEN technology an edge [25].

The use of self-organization as the dominant means of circuit assembly imposes the most severe limitation on nanoscale architectures: precise device alignment will be very difficult to achieve. Chemical self-assembly, as a stochastic process, will produce precise alignment of structures only rarely, and manipulation of single nanoscale structures to construct large scale circuits is impractical at best. Therefore, the most reliable way to make device connections will be to generate them at the intersection of two wires coated with the appropriate molecular devices. This alleviates the need for precise or end-to-end connections, as devices are incorporated into the circuit out of sheer topological necessity. However, connecting together more than two devices at a time, or devices with more than two terminals, will likely prove extremely difficult to do *en masse* at the nanometer scale. Three-terminal devices such as molecular analogs of transistors will not be practical, since ultimately the terminals of these devices will need to be connected in precisely defined patterns to produce properly functioning circuits. In practical, near-term nanocomputer designs, the active components will be two-terminal devices such as diodes, configurable switches, and molecular RTDs.

3 A Plausible Nanoarchitecture

Having enumerated the constraints imposed by CAEN, we present the highlights of a plausible architecture and fabrication process for constructing a molecular integrated circuit. The process, a hybrid of molecular self-assembly and lithographically produced aperiodic structure, is inexpensive yet creates structures out of nanoscale components. At the nanoscale, all subprocesses are self-directed; only at the micron scale do we allow deterministic operations. The process is hierarchical, proceeding from basic components (e.g. wires and switches), through self-assembled arrays of components, to complete systems. Since practical transistors will be difficult to achieve at the nanoscale, computation is performed using diode/resistor-based threshold logic.

In the first fabrication step, wires of different types are constructed through chemical self-assembly in bulk. Some of these wires may be functionalized, either by coating their exterior or by including inline devices along the axis of the wire. The next step aligns groups of wires into rafts of parallel wires using flow techniques [26]. Two rafts can be combined to form a two-dimensional grid using the same flow techniques. The active devices are created wherever two appropriately functionalized wires intersect. By choosing the molecules that coat the wires appropriately, we can make the intersection points ohmic contacts, diodes, configurable switches, or other devices. No precise alignment is required to create the devices; they occur wherever the rafts of wires intersect.

The resulting grids will be on the order of a few microns in size. A separate process will create a silicon-based die using standard lithography. Note, however, that this lithographic process does not have to be “state-of-the-art” since the CMOS components are few and far apart. The circuits on this die will provide power, clock lines, an I/O interface, and support logic for the grids of devices. The die will contain “holes” in which the grids are placed, aligned, and connected with the wires on the die.

To create useful aperiodic circuits using CAEN, reconfigurable devices will be incorporated into the grids. The main active component at the intersections of the grid will be a reconfigurable molecular diode (or a switch in series with a diode). The diode can be configured “off”, in which




Figure 1: A two-input AND gate implemented in a reconfigurable CAEN grid.

case the it acts like an open circuit, or “on,” in which case it acts like a diode. Figure 1 shows how a 3x3 grid might be configured to implement an AND gate. Four of the nine reconfigurable diodes are configured “on.” Since inverters cannot be constructed using diodes and resistors, CMOS inverters will be required at the input to the CAEN-based logic, which must compute both the desired output and its complement at every stage in the computation.

The nondeterministic nature of self-assembly will give rise to high defect densities. Instead of trying to eliminate these defects completely with manufacturing techniques, we perform post-fabrication steps that allow the chip to work in spite of its defects. A natural method of handling the defects, first used in the Teramac [27], is to exploit the reconfigurable nature of the device. First, the device is configured for self-diagnosis. The result of the self-diagnosis phase is a defect map. Then the defect map is used to configure the desired functionality around the defects. Reconfigurability is thus integral to the operation of a CAEN-based architecture.

4 Signal Restoration Without Transistors

The architecture described above satisfies the key constraints of nanoscale technology: it uses no deterministic operations at the nanoscale, and it uses only two-terminal nanoscale devices. To build an operational circuit, however, requires signal restoration between nanoscale wire grids. Connections to ground will remove current from the circuit, and every diode will cause a small voltage drop. CMOS transistors could be used between every grid to restore the signal to defined voltage and current levels. However, this is unattractive for two reasons. First, CMOS transistors are significantly larger than CAEN-based devices; requiring their use at every block would result in an overall decrease in the density of the fabric [28]. Second, if a low-current CAEN signal is forced to drive a CMOS signal-restoration transistor, the relatively large capacitance of the CMOS transistor will decrease the circuit’s overall speed.

Here we describe an alternative: a molecular latch, constructed entirely from molecular-scale devices, that can perform signal restoration using power from the clock to provide gain. This idea is motivated by early work on tunnel diodes [29] and more recently on compound semiconductor resonant tunneling diodes [30]. The molecular latch provides the important properties of signal restoration, I/O isolation, and noise immunity.




Figure 2: I - V curve for a “typical” RTD device.




Figure 3: (a) Basic molecular latch configuration. (b) Load-line diagram for molecular latch.

4.1 Molecular Latch Operation

The molecular latch is constructed from a pair of molecular resonant tunneling diodes (RTDs). Figure 2 shows an I - V curve of a representative RTD. The key feature of an RTD is a region of negative differential resistance (NDR), where the tunneling current falls as the voltage increases. The ratio of the current at the beginning of the NDR region to the current at the end is called the peak-to-valley ratio (PVR). Molecular RTDs that operate at room temperature with PVRs of 10 or more have already been realized [7, 31]

Figure 3 shows the arrangement of the RTDs in a molecular latch and a load-line diagram for the molecular latch at a voltage V_{ref} . The state of the latch is determined by the voltage at the node between the two RTDs. The load line diagram shows two stable states, V_{low} and V_{high} , and a third metastable state. Small voltage fluctuations in the metastable state will push the circuit into one or the other of the stable states. The low state represents binary “0” and the high state binary “1.”

The state of the latch is changed by temporarily disrupting and restoring this bistable equilibrium state. The bias voltage is temporarily lowered to V_{mono} , causing a shift of the load line to the left such that the circuit has only one stable state. The bias voltage is then returned to V_{ref} . During the evolution from the monostable to bistable state (the monostable-bistable transition, or MBT), current introduced to the data node will flow through the drive RTD. If the total charge introduced

by this current during the transition exceeds a threshold value, the circuit will switch to the high state; otherwise the circuit will switch to the low state. The amount of current necessary to force the latch into the high state is determined by the PVR of the RTD. For a well-matched RTD pair, this current can be very small [29]. A higher PVR requires more current to set the high state, but provides greater stability against current variation. For the circuit to function correctly, the drive RTD must be more resistive than the load RTD, i.e., R_{ser} in Figure 4a must be larger for RTD_{drive}.

RTD latch technology has already been exploited in constructing a GaAsFET logic family by Mathews et al. [30]. They describe a series of logic gates constructed from RTD latches, FETs, and saturated resistors that display high switching speeds and low power dissipation. In CAEN, the RTD latches are used for a different purpose: they provide voltage buffering by storing the state of previous computation for use in later computation, and signal restoration to either the '0' or '1' voltage with each latch. All computation is performed by diode-resistor logic and the latch restores the signal for a later computational stage. Work in the 1960s with tunnel diodes and threshold logic was hampered by high interdevice manufacturing variability [32]; the molecular RTDs discovered thus far have higher PVRs than semiconductor-based devices, which, as we explain below, may alleviate some of these problems.

4.2 Molecular Latch Simulations: Restoration and Reliability

Since molecular-scale RTD latches have yet to be realized, we conducted simulations of their behavior using SPICE. We used a similar device model as Mathews et al. for each RTD [30] The device model is depicted in Figure 4a. The RTD is modeled as a series resistance (R_{ser}), with a parallel capacitance (C_{RTD}) and voltage-dependent current source ($I(V)$). This meshes well with the proposed construction of molecular-scale RTDs, which is accomplished by connecting two segments of molecular-scale metal wire with a number of identical RTD molecules (See Figure 4b). Based on the estimated length and dielectric of the molecules and diameter of the molecular wires, we estimate a capacitance of 2.7 aF for each RTD. We constructed model equations for the voltage-dependent current source to mimic the I-V curves reported by Reed and Tour at 190K [9]. To explore the RTD parameter space, we also constructed models for several different "ideal" molecular RTDs in which the V_{peak} and V_{valley} locations and PVR were varied.

Figure 5(b) shows the results of a simple RTD simulation using an idealized molecule as the basis of the RTD latch circuit shown in Figure 5(a). The clock signal is shaped to provide a monostable-bistable transition (MBT) in every cycle and to maintain the bistable state for the remainder of the cycle. Note that the cycle time was deliberately made longer than required to ensure equilibrium before the next cycle. The figure shows that the voltage is restored to the levels predicted by the RTD load line and the choice of V_{ref} . In contrast to the "ideal" molecule, the relatively high V_{peak} of the Reed-Tour molecule would produce narrowly spaced '0' and '1' voltages for any reasonable choice of V_{ref} . The molecular RTD latch becomes more practical as V_{peak} becomes lower; the closer V_{peak} is to 0V, the larger the separation of low and high signals that can be produced at a realistically achievable V_{ref} .

To determine the envelope of stability for the RTD, we simulated the behavior of this simple circuit while varying the 'low' and 'high' input voltage, the input current, and the relative sizes of the load and drive RTDs. In simulation, simple RTD latches in isolation appear to be stable over the range of variability likely to be encountered in their synthesis. Stability has proved a significant




Figure 4: (a) Device model for RTD. (b) Possible construction technique for molecular RTD as an inline device (not to scale).




Figure 5: (a) Simple molecular latch circuit used in testing. (b) Results of single stability experiment. Note voltage “gain.” and memory effect.




Figure 6: *Simple 3-stage delay circuit using molecular latch from Figure 8.*




Figure 7: *Two-phase clocking scheme used for series latches.*

stumbling block to the application of semiconductor RTDs, as device variability often falls outside the required range for incorporation into circuits [32]. Molecular RTDs have an advantage in that all molecules incorporated into the RTD are identical. In addition, molecular RTDs have far bigger PVRs than conventional RTDs, which improves the stability of latches constructed with them [30].

4.3 Combinations of latches: I/O Isolation

In molecular circuits, latches are used to buffer and condition the output of a diode-resistor combinational circuit so that it can serve as input to the next one. We therefore need to consider the interactions of multiple latches. The simplest circuit using more than one latch is a delay circuit, as shown in Figure 6. Because the latch data node emits the correct voltage only when the latch clock is in the high state, a two-phase clocking scheme (Figure 7) is required to propagate the signal through the circuit. Latches connected in series have alternating clock phases.

It is clear that the lack of I/O isolation provided by transistors is problematic for a latch consisting solely of an RTD pair. Without intervening devices, all data nodes must exist at the same voltage, and even with additional linear circuit elements current is free to flow backwards to set up-stream latches. To provide the necessary isolation characteristics, we incorporate several additional devices into the latch, as shown in Figure 8.

A molecular-scale diode is added to prevent current from a downstream latch from flowing to




Figure 8: *The complete molecular latch.*




Figure 9: *AND, OR, XOR circuits using diode-resistor logic and intermediate molecular latches.*

set an upstream latch. Without the diode, setting a downstream latch will prevent the immediately upstream latch from resetting properly. Molecular diodes with very low voltage drops and reverse current flow have already been realized [33]. The resistor immediately following the latch data forces enough current into $\text{RTD}_{\text{drive}}$ to allow it to switch high, while allowing enough current into the next latch to determine its value. Without the resistor, latches in series would act as a current divider, and insufficient current would flow through each $\text{RTD}_{\text{drive}}$ to set it to the '1' state.

To understand the rationale for the resistor to ground, consider the delay latch shown in Figure 6 and the circuit state ($\text{In}=\text{Low}$, $\text{D1}=\text{Low}$, $\text{D2}=\text{High}$, $\text{D3}=\text{High}$) at the instant before the Φ_2 clock cycle. As the Φ_2 clock voltage is lowered to V_{mono} , the diodes in both D2 and D3 are reverse biased, preventing any current discharge from the capacitance of latch D2. Since discharge is necessary for the latch to enter the low state, we provide a path to ground through a large resistor. The value of this resistor must be large to prevent current loss that would cause latch-setting failure. However, the resistor should also be as small as possible in order to minimize the latch reset time.

We simulated several simple circuits using the above latch configuration, including the delay circuit in Figure 6 and the AND, OR, and XOR circuits shown in Figure 9. The circuits calculate the correct values, taking into account the delay required to traverse intervening latches. Circuits with more logic levels were also simulated successfully.

One additional complication resulting from a lack of I/O isolation was discovered for the AND




Figure 10: *Lack of I/O isolation in AND circuit provides path from V_{dd1} to ground.*




Figure 11: *Clocking scheme used for V_{dd} .*

circuit. Consider the circuit state ($A=\text{low}$, $B=\text{high}$, $Y=\text{high}$, $Z=\text{high}$) immediately before Φ_1 begins its cycle (Figure 10). The diode of Z is reverse biased, so the current in V_{dd1} will flow through the AND diodes into the latches and to ground, preventing them from being reset. The latch cannot be protected by a diode to prevent this, because doing so leaves no path to sink current from the pull-up voltage. Without a path to ground, the pull-up would cause Z to be set high incorrectly. To counteract this, we introduce a clocking scheme for the pull-up voltage as well, as shown in Figure Figure 11. The pull-up voltage is temporarily brought to ground during the MBT of the preceding latch. This removes the forward influence and allows the latches to be set properly. Unfortunately, in a real circuit this will likely have deleterious consequences for the RC constant.

4.4 Latch Summary

In simulation, the molecular latch is capable of restoring voltage levels to the two stable values and providing current input, as determined by the IV curve of the underlying RTD. It appears to be stable against variability brought about by manufacture, as long as the drive RTD can be reliably made more resistive than the load RTD. Because of the relatively low current required to switch states, one molecular latch can drive several other latches, i.e., latches can have moderate fan-out. In combination with the clocking scheme described above, the latch also provides the necessary I/O isolation to ensure proper calculation.

Some caveats are necessary, however. Because the RTD devices discovered thus far are ex-

tremely resistive, the RC constant for circuits incorporating them is very long (on the order of hundreds of nanoseconds). Clearly, this will need to be improved before these devices become practical. Also, the relative sizes of the various resistors must be carefully controlled during manufacture to ensure that they are properly matched; this will allow more devices to attain proper I/O isolation without adversely affecting the time constants for the circuit. Lastly, the use of a clocked pullup voltage may result in additional complications for the design of the CMOS support circuitry. However, the similarity in waveforms between the latch clock and V_{dd} clock may alleviate some of these problems.

5 Conclusions

Any architecture based on CAEN must overcome the constraints associated with directed assembly of nanometer-scale components and yet exploit the advantages of molecular electronics. Because the fabrication process is fundamentally nondeterministic and prone to introducing defects, the device must be reconfigurable and amenable to self-testing. This will allow the characteristics and defects of each device to be discovered, support the creation of circuits that avoid the defects, and fully exploit the available resources. For the foreseeable future, the fabrication processes will produce only simple, regular geometries. Therefore, proposed architectures should be built out of simple, two-dimensional, homogeneous structures. Rather than fabricating complex circuits, CAEN will rely on the reconfigurability of regular fabrics to implement arbitrary functions after manufacture. The construction process also needs to be parallel; heterogeneity should be introduced only at a lithographic scale. The resulting devices will be much like today's FPGAs in that they can be configured (and reconfigured) to implement any circuit in the field, but they will have several orders of magnitude more resources.

Electronic nanotechnology is quickly progressing and promises small, dense, low-power devices. Harnessing this power will require new ways of thinking about the manufacturing process. We will no longer be able to manufacture devices deterministically. Instead, post-fabrication reconfiguration will circumvent defects and determine the properties of the device.

References

- [1] R. W. Keyes. *Science* **230** 138–44 (1985).
- [2] R. Chau *et al.*, in *International Electron Devices Meeting* (December 2000).
- [3] Semiconductor Industry Association. *SIA Roadmap* (1997).
- [4] Sematech. *International Technology Roadmap for Semiconductors* (2000 Update).
- [5] P. Peery. *Nature* **406** 1023–1026 (2000).
- [6] J. Ellenbogen, J. Love. *Proc. IEEE* **88** 386–426 (2000).
- [7] J. Chen, M. A. Reed, A. M. Rawlett, J. M. Tour. *Science* **286** 1550–2 (1999).

-
- [8] J. Chen, M. Reed, in *Molecular Electronics 2000* (December 2000).
 - [9] J. Chen *et al.* *Appl. Phys. Lett.* **77** 1224–1226 (2000).
 - [10] C. Collier *et al.* *Science* **285** 391–3 (1999).
 - [11] N. S. M.V. Martinez-Diaz, J. Stoddart. *Ang. Chem. Intl. Ed. Eng.* **36** 1904 (1997).
 - [12] A. N. *et al.* *Appl. Phys. Lett* **70** 1742 (1997).
 - [13] T. Rueckes *et al.* *Science* **289** 94–97 (2000).
 - [14] S. J. Tans *et al.* *Nature* **386** 474–7 (1997).
 - [15] H. Park, A. Lim, J. Park, A. Alivisatos, P. McEuen. [www.physics.berkeley.edu/-research/mceuen/topics/nanocrystal/EMPaper.pdf](http://www.physics.berkeley.edu/research/mceuen/topics/nanocrystal/EMPaper.pdf) (1999).
 - [16] P. Collins, M. Arnold, P. Avouris. *Science* **292** 706–709 (2001).
 - [17] B. Martin *et al.* *Advanced Materials* **11** 1021–25 (1999).
 - [18] J. Mbindyo *et al.* *Advanced Materials* (2000).
 - [19] Xilinx Corporation. Virtex series fpgas. <http://www.xilinx.com/products/virtex.htm>.
 - [20] Altera Corporation. Apex device family. <http://www.altera.com/html/products/apex.html>.
 - [21] A. Ye, A. Moshovos, S. Hauck, P. Banerjee, in *Proceedings of the 27th Annual International Symposium on Computer Architecture* (June 2000).
 - [22] S. Goldstein *et al.*, in *Proceedings of the 26th Annual International Symposium on Computer Architecture*, 28–39 (May 1999).
 - [23] W. Culbertson, R. Amerson, R. Carter, P. Kuekes, G. Snider, in *Proc. IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems* (1996).
 - [24] For the CAEN device we assume that the nanowires are on 20nm centers. A CMOS transistor with a 4:1 ratio in a 70nm process (even using Silicon on Insulator, which does not need wells) with no wires attached measures 210nm x 280nm. Attaching minimally-sized wires to the terminals increases the size to 350nm x 350nm, which is still significantly smaller than the average size of a transistor in a real design.
 - [25] If we assume 20nm nanowire centers, then using the nanoblock architecture we describe and assuming 10 wires per nanoblock (7 internal and 3 external), one nanoblock is 200x200nm = 40,000nm², excluding Vdd/Gnd and latches. Even if we assume that these elements increase the area tenfold to 400,000nm², the nanoblock can implement a full adder in the area of one CMOS transistor.
 - [26] Y. Huang, X. Duan, Q. Wei, C. Lieber. *Science* **291** 630–3 (2001).

-
- [27] W. Culbertson, R. Amerson, R. Carter, P. Kuekes, G. Snider, in *Proceedings of the 1997 IEEE Symposium on FPGAs for Custom Computing Machines*, 116–124 (April 1997).
- [28] S. Goldstein, M. Budiu, in *Proceedings of the 28th Annual International Symposium on Computer Architecture* (June 2001).
- [29] I. Aleksander, R. Scarr. *J British IRE* **23** 177–192 (1962).
- [30] R. Mathews *et al.* *Proc. IEEE* **87** 596 (1999).
- [31] J. Chen, M. Reed, in *Proceedings of International Conference on Molecular Electronics*, Kona, HI (Dec. 2000).
- [32] D. Holmes, P. Baynton, in A. Strickland, editor, *Gallium Arsenide*, 3, 236–240, London (1967). Institute of Physics and the Physical Society.
- [33] N. I. Kovtyukhova *et al.* Submitted to *J. Phys. Chem.*
- [34] We would like to thank Theresa Mayer, Herman Schmidt, and Krste Asanovic for their insights and constructive criticism of the drafts of this work. This work was supported in part by DARPA, an NSF CAREER grant, and Hewlett-Packard Corporation.